

Learning in embodied action-perception loops through exploration

Daniel Y. Little and Friedrich T. Sommer

Abstract—Although exploratory behaviors are ubiquitous in the animal kingdom, their computational underpinnings are still largely unknown. Behavioral Psychology has identified learning as a primary drive underlying many exploratory behaviors. Exploration is seen as a means for an animal to gather sensory data useful for reducing its ignorance about the environment. While related problems have been addressed in Data Mining and Reinforcement Learning, the computational modeling of learning-driven exploration by embodied agents is largely unrepresented.

Here, we propose a computational theory for learning-driven exploration based on the concept of missing information that allows an agent to identify informative actions using Bayesian inference. We demonstrate that when embodiment constraints are high, agents must actively coordinate their actions to learn efficiently. Compared to earlier approaches, our exploration policy yields more efficient learning across a range of worlds with diverse structures. The improved learning in turn affords greater success in general tasks including navigation and reward gathering. We conclude by discussing how the proposed theory relates to previous information-theoretic objectives of behavior, such as predictive information and the free energy principle, and how it might contribute to a general theory of exploratory behavior.

Index Terms—Knowledge acquisition, Information theory, Control theory, Machine learning, Psychology, Computational neuroscience.



1 INTRODUCTION

EXPLORATORY behaviors have been observed and studied in diverse species across the animal kingdom. As one example, approach and investigation of novel stimuli have been studied in vertebrates ranging from fish to birds, reptiles, and mammals [21], [40], [41], [60], [68], [75]. As another, open field and maze experimental paradigms for studying locomotive exploration in mice and rats have recently been adapted to behavioral studies in zebrafish [64], [65]. Indeed, exploratory behaviors have even been described across a range of invertebrates [12], [26], [31], [52]. The prevalence of exploratory behaviors across animal species suggests a fundamental evolutionary advantage, largely believed to derive from the utility of information acquired through such behaviors [33], [50], [51], [55], [56].

Computational models of exploratory behavior, developed predominantly in the field of Reinforcement Learning (RL), have largely focused on the role of exploration in the acquisition of external rewards [1], [7], [32], [34], [69], [70], [71]. An agent that strictly maximizes the acquisition of known rewards might fall short in finding new, previously unknown, sources of reward. Reward maximization therefore requires balancing between directed harvesting of known rewards (exploitation) and

the search for new rewards (exploration) [1], [7], [32], [34], [58], [71]. The emphasis in the RL literature on reward acquisition, however, stands in contrast to the dominant psychological theories of exploration. Quoting D. E. Berlyne, a pioneer in the psychology of exploration:

As knowledge accumulated about the conditions that govern exploratory behavior and about how quickly it appears after birth, it seemed less and less likely that this behavior could be derivative of hunger, thirst, sexual appetite, pain, fear of pain, and the like, or that stimuli sought through exploration are welcomed because they have previously accompanied satisfaction of these drives. [11]

Berlyne further suggested “the most acute motivational problems . . . are those in which the perceptual and intellectual activities are engaged in for their own sake and not simply as aids to handle practical problems” [10]. In fact, a consensus has emerged in behavioral psychology that learning represents the primary drive of exploratory behaviors [2], [37], [52], [62]. To address this gap between computational modeling and behavioral psychology, we introduce here a mathematical framework for studying how behavior effects learning and develop a novel model of learning-driven exploration.

In Computational Neuroscience, machine learning techniques have been successfully applied towards modeling how the brain might learn the structure underlying sensory signals, e.g., [15], [16], [18], [36], [47], [54], [59]. Generally, these methods focus on passive learning where the learning system can not directly effect the

- D.Y. Little is with the Department of Molecular and Cellular Biology, University of California at Berkeley, Berkeley, CA 94720.
E-mail: dylittle@berkeley.edu
- F.T. Sommer is with the Redwood Center for Theoretical Neuroscience, University of California, Berkeley, CA 94720
E-mail: fsommer@berkeley.edu

sensory input it receives. Exploration, in contrast, is inherently active, and can only occur in the context of a closed-action perception loop. Learning in closed action-perception loops differs from passive learning in two important aspects [22]. First, a learning agent’s internal model of the world must keep track of how actions change the sensory input. Sensorimotor contingencies, such as the way visual scenes change as we shift our gaze or move our head, must be taken into account to properly attribute changes in sensory signals to their causes. This is perhaps reflected in neuroanatomy where tight sensory-motor integration has been reported at all levels of the brain [23], [24]. Though often taken for granted, sensor-motor contingencies must actually be learned during the course of development as is eloquently expressed in the explorative behaviors of young infants (e.g., grasping and manipulating objects during proprioceptive exploration or bringing them into visual view during intermodal exploration) [45], [48], [57].

The second crucial aspect of learning in a closed action-perception loop is that actions direct the acquisition of sensory data. To discover what is inside an unfamiliar box, a curious child must open it. To learn about the world, scientists perform experiments. Directing the acquisition of data is particularly important for embodied agents whose actuators and sensors are physically confined. Since the most informative data may not always be accessible to a physical sensor, embodiment may constrain an exploring agent and require that it coordinates its actions to retrieve useful data.

In the model we propose here, an agent moving between discrete states in a world has to learn how its actions influence its state transitions. The underlying transition dynamics are governed by a Controllable Markov Chain (CMC). Within this simple framework, various utility functions for guiding exploratory behaviors will be studied, as well as several methods for coordinating actions over time. The different exploratory strategies are compared in their rate of learning.

2 MODEL

2.1 Mathematical framework for embodied active learning

Controllable Markov chains (CMCs) are a simple extension of Markov chains that incorporate a control variable for switching between different transition distributions [19]. Formally, a CMC is a 3-tuple $(\mathcal{S}, \mathcal{A}, \Theta)$ where:

- \mathcal{S} is a finite set of the possible *states* of the system (for example, the possible location of an agent in its world). $N = |\mathcal{S}|$
- \mathcal{A} is a finite set of the possible control values, i.e., the *actions* the agent can choose. $M = |\mathcal{A}|$
- Θ is a 3-dimensional CMC *kernel* describing the probability of transitions between states given an action (for example, the probability an agent moves

from state s to state s' when it chooses action a):

$$p_{a,s}(s'|\Theta) = \Theta_{a,s,s'} \quad (1)$$

$$\sum_{s'} \Theta_{a,s,s'} = 1$$

For any fixed action, $\Theta_{a,:}$ is a (two-dimensional) stochastic matrix describing a Markov process. Each column in this matrix $\Theta_{a,s,:}$ defines a *transition distribution* which is a categorical (finite discrete) distribution specifying the likelihoods for the next state.

The CMC provides a simple mathematical framework for modeling exploration in embodied action-perception loops. At every time step, an exploring agent is allowed to freely select any action $a \in \mathcal{A}$. The learning task of the exploring agent is to build from observed transitions an estimate, the *internal model* $\hat{\Theta}$, of the true CMC kernel, the *world* Θ . We assume that the explorer begins with limited information about the world in the form of a prior and must improve its estimate by acting in the world and gathering observations. The states can be directly observed by the agent, i.e. the system is not hidden. In the CMC framework, an agent’s immediate ability to interact with and observe the world is limited by the current state. This restriction models the *embodiment* of the agent. To ameliorate the myopia imparted by its embodiment, an agent can coordinate its actions over time. Our primary question is how action policies can optimize the speed and efficiency of learning in embodied action-perception loops.

2.2 Information-theoretic assessment of learning

To assess an agent’s success towards learning, we determine the *missing information* I_M in its internal model (as proposed by Pfaffelhuber [49]). We do this by first calculating the Kullback-Leibler (KL) divergence of the internal model from the world for each transition distribution:

$$D_{KL}(\Theta_{a,s,:} \parallel \hat{\Theta}_{a,s,:}) := \sum_{s'=1}^N \Theta_{a,s,s'} \log_2 \left(\frac{\Theta_{a,s,s'}}{\hat{\Theta}_{a,s,s'}} \right)$$

The KL-divergence can be interpreted as the expected amount of information in bits, lost when observations (following the true distribution) are communicated using an encoding scheme optimized for the estimated distribution [13]. The loss is large when the two distributions differ greatly and zero when they are identical. The missing information is then calculated as the sum of the KL-divergences for each transition distribution:

$$I_M(\Theta \parallel \hat{\Theta}) := \sum_{s \in \mathcal{S}, a \in \mathcal{A}} D_{KL}(\Theta_{a,s,:} \parallel \hat{\Theta}_{a,s,:}) \quad (2)$$

We will use missing information (2) to assess learning and to compare the performance of different explorative strategies. Steeper decreases in missing information over time represent faster learning and thus more efficient explorative strategies. Table 1 has been included as a reference guide for the various measures discussed in this manuscript for assessing or guiding exploration.

TABLE 1
Table of Information Measures

Name used here, Abbreviation (Equation Number)	Name used in [Reference]	Mathematical expression
<i>Missing Information</i> , I_M (2)	<i>Missing Information</i> [49]	$\sum_{s \in \mathcal{S}, a \in \mathcal{A}} D_{KL}(\Theta_{s,a,:} \hat{\Theta}_{s,a,:})$
<i>Information Gain</i> , I_G (10)		$I_M(\hat{\Theta}) - I_M(\hat{\Theta}^{a,s,s^*})$
<i>Predicted Information Gain</i> , PIG (13)	<i>Information Gain</i> [44]	$E_{s^* a,s,\vec{d}_{a,s}} \left[D_{KL}(\hat{\Theta}_{a,s,:} \hat{\Theta}_{a,s,:}^{a,s,s^*}) \right]$
<i>Posterior Expected Information Gain</i> , $PEIG$ (18)	<i>KL – Divergence</i> [67] or <i>Surprise</i> [29]	$\sum_{s \in \mathcal{S}, a \in \mathcal{A}} D_{KL}(\hat{\Theta}_{a,s,:}^{past} \hat{\Theta}_{a,s,:}^{current})$
<i>Predicted Mode Change</i> , PMC (19)	<i>Probability Gain</i> [44]	$\sum_{s^*} \hat{\Theta}_{a,s,s^*} \left[\max_{s'} \hat{\Theta}_{a,s,s'}^{a,s,s^*} - \max_{s'} \hat{\Theta}_{a,s,s'} \right]$
<i>Predicted L_1 Change</i> , PLC (20)	<i>Impact</i> [44]	$\sum_{s^*} \hat{\Theta}_{a,s,s^*} \left[\frac{1}{N} \sum_{s'} \left \hat{\Theta}_{a,s,s'}^{a,s,s^*} - \hat{\Theta}_{a,s,s'} \right \right]$

2.3 Bayesian inference learning in an agent

During exploration, a learning agent will gather data in the form of experienced state transitions. At every time step, it can update its internal model $\hat{\Theta}$ with its last observation.

Assuming the transition probabilities are drawn from a *prior distribution* f (a distribution over CMC kernels), and letting \vec{d} be the history of observed transitions, the Bayesian estimate $\hat{\Theta}_{a,s,s'}$ for the probability of transitioning to state s' from state s under action a is given by:

$$\begin{aligned}
\hat{\Theta}_{a,s,s'} &:= p_{a,s}(s'|\vec{d}) \\
&= \int_{\Theta} p_{a,s}(s', \Theta|\vec{d}) d\Theta \\
&= \int_{\Theta} p_{a,s}(s'|\Theta, \vec{d}) f(\Theta|\vec{d}) d\Theta \\
&= \int_{\Theta} p_{a,s}(s'|\Theta) f(\Theta|\vec{d}) d\Theta \\
&= \int_{\Theta} \Theta_{a,s,s'} f(\Theta|\vec{d}) d\Theta \\
&= E_{\Theta}[\Theta_{a,s,s'}|\vec{d}] := E_{\Theta|\vec{d}}[\Theta_{a,s,s'}] \quad (3)
\end{aligned}$$

For discrete priors the above integrals would be replaced with summations. At the beginning of exploration, when the history \vec{d} is empty, $f(\Theta|\vec{d})$ is simply the prior $f(\Theta)$. Equation (3) demonstrates that the Bayesian estimate is equivalent to the expected value of the true CMC kernel given the data. If we further assume that each transition distribution $\Theta_{a,s,:}$ is independently drawn from the marginals $f_{a,s}$ of the distribution f and we let $\vec{d}_{a,s}$ be the history of state transitions experienced when taking action a in state s , (3) simplifies to the

independent estimation of each transition distribution:

$$\begin{aligned}
\hat{\Theta}_{a,s,s'} &= \int_{\Theta} \Theta_{a,s,s'} f(\Theta|\vec{d}) d\Theta \\
&= \int_{\Theta_{a,s,:}} \Theta_{a,s,s'} f_{a,s}(\Theta_{a,s,:}|\vec{d}_{a,s}) d\Theta_{a,s,:} \\
&= E_{\Theta_{a,s,:}|\vec{d}_{a,s}}[\Theta_{a,s,s'}] \quad (4)
\end{aligned}$$

In the following theorem, we demonstrate that the Bayesian estimator minimizes the expected missing information and thus is the best estimate under our objective function (2).

Theorem 1: Let Θ and Φ be CMC kernels. Θ describes the ground truth environment generated from a prior distribution. Φ is any internal model of the agent. Then the expected missing information between Θ and an internal model Φ , given data \vec{d} , is minimized by the Bayesian estimate $\hat{\Theta}$:

$$\hat{\Theta} = \arg \min_{\Phi} E_{\Theta|\vec{d}}[I_M(\Theta || \Phi)] \quad (5)$$

Proof: Since missing information is simply the sum of the KL-divergence for each transition kernel (2), minimizing missing information is equivalent to independently minimizing these KL-divergences:

$$\begin{aligned}
&\arg \min_{\Phi_{a,s,:}} E_{\Theta_{a,s,:}|\vec{d}}[D_{KL}(\Theta_{a,s,:} || \Phi_{a,s,:})] \\
&= \arg \min_{\Phi_{a,s,:}} E_{\Theta_{a,s,:}|\vec{d}} \left[\sum_{s'} \Theta_{a,s,s'} \log_2 \left(\frac{\Theta_{a,s,s'}}{\Phi_{a,s,s'}} \right) \right] \\
&= \arg \min_{\Phi_{a,s,:}} E_{\Theta_{a,s,:}|\vec{d}} \left[\sum_{s'} \Theta_{a,s,s'} \log_2 \Theta_{a,s,s'} \right] \\
&\quad - E_{\Theta_{a,s,:}|\vec{d}} \left[\sum_{s'} \Theta_{a,s,s'} \log_2 \Phi_{a,s,s'} \right]
\end{aligned}$$




Fig. 1. Example maze. The 36 states correspond to rooms in a maze. The 4 actions correspond to noisy translations in the cardinal directions. Two transition distributions are depicted, each by a set of 4 arrows emanating from their starting state. Flat-headed arrows represent translations into walls, resulting in staying in the same room. Dashed arrows represent translation into a portal (blue lines) leading to the base state (blue target). The shading of an arrow indicates the probability of the transition (darker color represents higher probability).

$$\begin{aligned}
 &= \arg \min_{\Phi_{a,s,:}} -E_{\Theta_{a,s,:}|\vec{d}} \left[\sum_{s'} \Theta_{a,s,s'} \log_2 \Phi_{a,s,s'} \right] \\
 &= \arg \min_{\Phi_{a,s,:}} - \sum_{s'} E_{\Theta_{a,s,:}|\vec{d}} [\Theta_{a,s,s'}] \log_2 \Phi_{a,s,s'} \\
 &= \arg \min_{\Phi_{a,s,:}} H \left[E_{\Theta_{a,s,:}|\vec{d}} [\Theta_{a,s,:}] ; \Phi_{a,s,:} \right]
 \end{aligned}$$

Here $H[\theta; \phi]$ denotes the cross-entropy [13]. Then, by Gibb's inequality [13] we conclude:

$$\begin{aligned}
 &\arg \min_{\Phi_{a,s,:}} H \left[E_{\Theta_{a,s,:}|\vec{d}} [\Theta_{a,s,:}] ; \Phi_{a,s,:} \right] \\
 &= E_{\Theta_{a,s,:}|\vec{d}} [\Theta_{a,s,:}] \\
 &= \hat{\Theta}_{a,s,:}(\vec{d})
 \end{aligned}$$

□

The analytical form for the Bayesian estimate will depend on the prior. In the following section, we introduce three classes of CMCs which will be considered in this study and specify the Bayesian estimates for each.

2.4 Three test environments for studying exploration

In the course of exploration, the data an agent accumulates will depend on both its behavioral strategy as well as the world structure. Studying diverse environments, i.e., CMCs that differ greatly in structure, will help us to investigate how world structure effects the relative performance of different exploratory strategies and to identify action policies that produce efficient learning under broad conditions.

The three different classes of test environments to be investigated will be called Dense Worlds, Mazes, and 1-2-3 Worlds. For each class, random CMCs are generated by drawing the transition distributions from specific distributions. These generative distributions are also given to the agents as priors for performing Bayesian inference.

1) *Dense Worlds* correspond to complete directed probability graphs with $N = 10$ states and $M = 4$ actions. Each transition distribution is independently drawn from a Dirichlet distribution:

$$\Theta_{s,a,:} \sim \text{Dir}(1) := \frac{\prod_{s'} \Theta_{s,a,s'}^{\alpha-1}}{B(\alpha)} \quad (6)$$

$$\begin{aligned}
 B(\alpha) &:= \frac{\Gamma(\alpha)^N}{\Gamma(N\alpha)} \\
 \Gamma(\alpha) &:= \int_0^\infty t^{\alpha-1} e^{-t} dt
 \end{aligned}$$

The Dirichlet distribution is the conjugate prior of the categorical distribution (1) and thus a natural distribution for generating CMCs. It is parametrized by a concentration factor α that determines how much weight in the Dirichlet distribution is centered at the midpoint of the simplex, the space of all possible transition distributions. The midpoint corresponds to the uniform categorical distribution. For Dense Worlds, we use a concentration parameter $\alpha = 1$ which results in a uniform distribution over the simplex. An example is depicted graphically in Fig. S1.

The Bayesian estimate for Dense Worlds has the following closed-form expression:

$$\hat{\Theta}_{a,s,s'} = \frac{1 + \sum_{k=1}^K \delta_{s', \vec{d}_{a,s}[k]}}{K + N} \quad (7)$$

where K is the length of the history $\vec{d}_{a,s}$ and $\delta_{x,y}$ is the Kronecker delta (i.e. $\delta_{x,y}$ is 1 if $x = y$ and 0 otherwise). Equation (7) reveals that the Bayesian estimate is simply the relative frequencies of the observed data with the addition of one fictitious count per transition. The incorporation of this fictitious observation is referred to as Laplace smoothing and is often performed to avoid over-fitting [39]. The derivation of Laplace smoothing from Bayesian inference over a Dirichlet prior is a well known result [38].

2) *Mazes* consist of $N = 36$ states corresponding to rooms in a randomly generated 6 by 6 maze and $M = 4$ actions corresponding to noisy translations, each biased towards one of the four cardinal directions "up", "down", "left" and "right". An example is depicted in Fig. 1. Walking into a wall causes the agent to remain in its current location. There are 30 transporters randomly distributed amongst the walls which lead to a base state. Each maze has a single, randomly chosen base state (concentric rings in Fig. 1). All transitions that do not correspond to a single translation are assigned a probability of zero. The non-zero probabilities are drawn from a Dirichlet distribution with concentration parameter $\alpha = 0.25$. The highest probability is assigned

to the state corresponding to the cardinal direction of the action. The small concentration parameter distributes more probability weight in the corners of the simplex corresponding to deterministic transitions. This results in Maze transitions have strong biases towards an actions associated cardinal direction.

Agents in Mazes must estimate the non-zero transitions using the Dirichlet prior without knowledge of each action's assigned cardinal direction. Similar to (7), the Bayesian estimate for maze transitions is given by:

$$\hat{\Theta}_{a,s,s'} = \frac{0.25 + \sum_{k=1}^K \delta_{s', \vec{d}_{a,s}[k]} }{K + 0.25 \cdot N_{a,s}} \quad (8)$$

where $N_{a,s}$ is the number of non-zero probability states in the transition distribution $\Theta_{a,s,:}$. As with Dense Worlds, the Bayesian estimate (8) for mazes is a Laplace smoothed histogram.

3) 1-2-3 Worlds consists of $N = 20$ states and $M = 3$ actions. In a given state, action $a = 1$ moves the agent deterministically to a single target state, action $a = 2$ brings the agent with probability 0.5 to one of two possible target states, and action $a = 3$ brings the agent with probability 0.333 to one of 3 potential target states. The target states are randomly and independently selected for each action taken in each state. To create an absorbing state, the probability that state 1 is among the targets of action a is set to $1 - 0.75^a$. The probability for all other states to be selected as targets is uniform. Explicitly, letting Ω_a be the set of all admissible transition distributions for action a :

$$\Omega_a := \{\theta \in \mathbb{R}^N \mid \sum_{s'} \theta_{s'} = 1 \text{ and } \theta_{s'} \in \{0, \frac{1}{a}\} \forall s'\}$$

the transition distributions are drawn from the following distribution:

$$p(\Theta_{a,s,:}) = \begin{cases} 0 & \text{if } \Theta_{a,s,:} \notin \Omega_a \\ \frac{1 - 0.75^a}{\binom{N-1}{a-1}} & \text{else if } \Theta_{a,s,1} = \frac{1}{a} \\ \frac{1 - (1 - 0.75^a)}{\binom{N-1}{a}} & \text{otherwise} \end{cases} \quad (9)$$

If this process results in a non ergodic CMC, it is discarded and a new CMC is generated. A CMC as ergodic if, for every ordered pair of states, there exist an action policy under which an agent starting at the first state will eventually reach the second. An example 1-2-3 Worlds is depicted in Fig. S2.

Bayesian inference in 1-2-3 Worlds differs greatly from Mazes and Dense Worlds because of its discrete prior. Essentially, state transitions that have been observed are accurately estimated, while the remaining probability weight is distributed across those states that have not yet been experienced (preferentially to state 1 and uniformly across other states). Explicitly, if $a, s \rightarrow s'$ has been previously observed, then the Bayesian estimate for $\Theta_{a,s,s'}$

is given by:

$$\hat{\Theta}_{a,s,s'} = \frac{1}{a}$$

If $a, s \rightarrow s'$ has not been observed but $a, s \rightarrow 1$ has, then the Bayesian estimate is given by:

$$\hat{\Theta}_{a,s,s'} = \frac{1 - \frac{|S^*|}{a}}{N - T}$$

Here T is the number of target states that have already been observed:

$$T := |\{s^* \in \vec{d}_{a,s}\}|$$

Finally, if neither $a, s \rightarrow s'$ nor $a, s \rightarrow 1$ have been observed, then the Bayesian estimate is:

$$\hat{\Theta}_{a,s,s'} = \begin{cases} \frac{1 - 0.75^a}{1 + \left(\binom{a-1}{T} - 1\right) * 0.75^a} \cdot \frac{1}{a} & \text{if } s' = 1 \\ \frac{1 - \left(\frac{T}{a} + \hat{\Theta}_{a,s,1}\right)}{N - T - 1} & \text{otherwise} \end{cases}$$

3 RESULTS

3.1 Assessing the information-theoretic value of planned actions

The central question to be addressed is how actions effect the learning process in embodied action-perception loops. Ideally, actions should be chosen so that the missing information (2) decreases as fast as possible. As discussed in Section 2.3, the Bayesian estimate minimizes the expected missing information. We will assume that an agent continually updates its internal model accordingly from the observations it receives. The Bayesian estimate, however, does not indicate which action will optimize the utility of future data. Towards this objective, an agent should try to predict the impact a new observation will have on its missing information. We call the decrease in missing information between two internal models the *information gain* (I_G). Letting $\hat{\Theta}$ be a current model derived from data \vec{d} and $\hat{\Theta}^{a,s,s^*}$ be an updated model derived from adding an observation of $a, s \rightarrow s^*$ to \vec{d} , the information gain for this observation is:

$$\begin{aligned} I_G(a, s, s^*) &:= I_M(\Theta \parallel \hat{\Theta}) - I_M(\Theta \parallel \hat{\Theta}^{a,s,s^*}) \\ &= D_{KL}(\Theta_{a,s,:} \parallel \hat{\Theta}_{a,s,:}) - D_{KL}(\Theta_{a,s,:} \parallel \hat{\Theta}_{a,s,:}^{a,s,s^*}) \\ &= \sum_{s'} \Theta_{a,s,s'} \log_2 \frac{\Theta_{a,s,s'}}{\hat{\Theta}_{a,s,s'}} - \Theta_{a,s,s'} \log_2 \frac{\Theta_{a,s,s'}}{\hat{\Theta}_{a,s,s'}^{a,s,s^*}} \\ &= \sum_{s'} \Theta_{a,s,s'} \log_2 \frac{\hat{\Theta}_{a,s,s^*}^{a,s,s^*}}{\hat{\Theta}_{a,s,s'}} \end{aligned} \quad (10)$$

Calculating the information gained from taking action a in state s would therefore require knowing Θ as well as s^* . An agent can only infer former and can only know the latter after it has executed the action. In the following theorem, however, we derive a closed-form expression




Fig. 2. Accuracy of predicted information gain. The average predicted information gain is plotted against the average realized information gain. Averages are taken over 200 CMCs, $N \times M$ transition distributions, and 50 trials. Error bars depict standard deviations (only plotted above the mean for 1-2-3 Worlds). The arrow indicates the direction of increasing numbers of observations (top-right = none, bottom-left = 19). The unity lines are drawn in gray.

for the expected information gain, which we shall call the *predicted information gain* (PIG). □

Theorem 2: Let Θ be a CMC kernel whose transition distributions are independently generated from prior distributions. If an agent is in state s and has previously collected data \vec{d} , then the expected information gain for taking action a and observing the resultant state s^* is given by:

$$\begin{aligned} \text{PIG}(a, s) &:= \mathbb{E}_{s^*, \Theta | \vec{d}} [\text{IG}(a, s, s^*)] \\ &= \sum_{s^*} \hat{\Theta}_{a, s, s^*} \text{D}_{\text{KL}}(\hat{\Theta}_{a, s, :} \| \hat{\Theta}_{a, s, s^*}^{a, s, s^*}) \end{aligned} \quad (11)$$

where $\hat{\Theta}$ is the current internal model of the agent and $\hat{\Theta}_{a, s, s^*}^{a, s, s^*}$ is what the internal model would become if it were updated with an observation s^* resulting from a prospective new action a .

Proof:

$$\begin{aligned} &\mathbb{E}_{s^*, \Theta | \vec{d}} [\text{IG}(a, s, s^*)] \\ &= \mathbb{E}_{s^*, \Theta | \vec{d}} \left[\sum_{s'} \Theta_{a, s, s'} \log 2 \left(\frac{\hat{\Theta}_{a, s, s'}^{a, s, s^*}}{\hat{\Theta}_{a, s, s'}} \right) \right] \\ &= \mathbb{E}_{s^*, \Theta_{a, s, :} | \vec{d}_{a, s}} \left[\sum_{s'} \Theta_{a, s, s'} \log 2 \left(\frac{\hat{\Theta}_{a, s, s'}^{a, s, s^*}}{\hat{\Theta}_{a, s, s'}} \right) \right] \\ &= \mathbb{E}_{s^* | \vec{d}_{a, s}} \left[\mathbb{E}_{\Theta | \vec{d}_{a, s}, s^*} \left[\sum_{s'} \Theta_{a, s, s'} \log 2 \left(\frac{\hat{\Theta}_{a, s, s'}^{a, s, s^*}}{\hat{\Theta}_{a, s, s'}} \right) \right] \right] \\ &= \mathbb{E}_{s^* | \vec{d}_{a, s}} \left[\sum_{s'} \mathbb{E}_{\Theta | \vec{d}_{a, s}, s^*} [\Theta_{a, s, s'}] \log 2 \left(\frac{\hat{\Theta}_{a, s, s'}^{a, s, s^*}}{\hat{\Theta}_{a, s, s'}} \right) \right] \\ &= \mathbb{E}_{s^* | \vec{d}_{a, s}} \left[\sum_{s'} \hat{\Theta}_{a, s, s'} \log 2 \left(\frac{\hat{\Theta}_{a, s, s'}^{a, s, s^*}}{\hat{\Theta}_{a, s, s'}} \right) \right] \\ &= \mathbb{E}_{s^* | \vec{d}_{a, s}} [\text{D}_{\text{KL}}(\hat{\Theta}_{a, s, :} \| \hat{\Theta}_{a, s, s^*}^{a, s, s^*})] \\ &= \sum_{s^*} p(s^* | \vec{d}_{a, s}) \text{D}_{\text{KL}}(\hat{\Theta}_{a, s, :} \| \hat{\Theta}_{a, s, s^*}^{a, s, s^*}) \\ &= \sum_{s^*} \hat{\Theta}_{a, s, s^*} \text{D}_{\text{KL}}(\hat{\Theta}_{a, s, :} \| \hat{\Theta}_{a, s, s^*}^{a, s, s^*}) \end{aligned}$$

Notice, (11) can be computed from previously collected data alone. For each class of environments, Fig. 2 compares the average PIG with the average realized information gain as successive observations are drawn from a transition distribution and used to update a Bayesian estimate. In accordance with Theorem 2, in all three environments PIG accurately predicts the average information gain. Thus, theoretically and empirically, PIG represents an accurate estimate of the average gains towards the learning objective functions that an agent can expect to receive for taking a planned action in a particular state.

Interestingly, the equation for computing PIG, RHS of (11), has been previously considered in the field of psychology, where it was applied to describe human behavior during hypothesis testing [35], [44], [46]. To our knowledge, however, its equality to the expected decrease in missing information (Theorem 2) has not been previously shown.

3.2 Control learners: unembodied and random action

During exploration, an embodied agent can choose its action but is bound to the state that resulted from its last transition. A simple exploratory strategy would be to always select actions uniformly randomly. We will use such a *random action* strategy as a baseline control for learning performance representing a naive explorer.

In contrast to embodied agents, one can also consider an *unembodied* agent that is allowed to arbitrarily relocate to a new state before taking an action. For an unembodied agent, optimization of learning becomes much simpler as it decomposes into an independent sampling problem [49]. Since the PIG for each transition distribution decreases monotonically over successive observations (Fig. 2), learning by an unembodied agent can be optimized by always sampling from the state and action pair with the highest PIG. Thus, learning can be




Fig. 3. Learning curves for control strategies. The average missing information is plotted over exploration time for the unembodied positive control and random action baseline control. Standard errors are plotted as dotted lines above and below learning curves. ($n=200$)

optimized in a greedy fashion:

$$(a, s)_{\text{Unemb.}} := \arg \max_{(a, s)} \text{PIG}(a, s) \quad (12)$$

The learning curves of the unembodied agent will serve here as a *positive control* as it represents an upper bound for the performance of embodied agents.

An initial comparison between random action and the unembodied control highlights a notable difference among the three classes of environments (Fig. 3). Specifically, the performance margin between the two controls is significant in Mazes and 1-2-3 Worlds ($p < 0.001$), but not in Dense Worlds ($p > 0.01$). The significance was assessed by post-hoc analysis of Friedman’s test [25] comparing the areas under the two learning curves. Despite using a naive strategy, the random actor is essentially reaching maximum performance in Dense Worlds, suggesting that exploration of this environment is fairly easy. The difference in performance between random action and the unembodied control offers an initial insight into the constraints experienced by embodied agents. A directed exploration strategy may help bridge this gap.

3.3 Exploration strategies based on PIG

Given that PIG can be computed by an agent using only the data it has already collected (along with its prior), we wondered whether it could be used as a utility function to guide exploration. Since greedy maximization of PIG is optimal for the unembodied agent, one might consider a similar greedy strategy for an embodied agent (PIG(greedy)). The key difference would be that the embodied agent can only select its action but not its current state:

$$a_{\text{PIG(greedy)}} := \arg \max_a \text{PIG}(a, s) \quad (13)$$

The performance comparison between PIG(greedy) (13) and the unembodied control (12) is of particular interest because the two strategies differ only in that one

is embodied but the other is not. Thus differences of their performance reflect the embodiment constraint on learning. As shown in Fig. 4 the performance difference is largest in Maze worlds, moderate though significant in 1-2-3 Worlds and smallest in Dense Worlds ($p < 0.001$ for Mazes and 1-2-3 Worlds, $p > 0.001$ for Dense Worlds). To quantify the embodiment constraint faced in a world, we define an *embodiment index* as the relative difference between the areas under the learning curves for PIG(greedy) and the unembodied control which average 0.02 for Dense Worlds, 2.59 for Mazes, and 1.27 for 1-2-3 Worlds.

Also of particular interest, the comparison between PIG(greedy) and random action provides further insight differentiating the three classes of worlds (Fig. 4). Whereas PIG(greedy) yielded no improvement over random action in Dense Worlds and Mazes ($p > 0.001$), it significantly improved learning in 1-2-3 Worlds ($p < 0.001$), demonstrating that agents benefitted from the information-theoretic utility function only in 1-2-3 Worlds.

Greedy maximization of PIG considers only the immediate gains available and fails to account for the effect an action can have on future utility. In particular, when the potential for information gain is unevenly distributed, it may be necessary to coordinate actions over time to obtain remote but informative observations. Forward estimation of total future PIG over multiple time steps is intractable as the number of action sequences and state outcomes increases exponentially with time. To guide an agent towards maximizing long-term gains of PIG, we instead employ a back-propagation approach previously developed in the field of economics, Value-Iteration (VI) [8]. The estimation starts at a distant time point (initialized as $\tau = 0$) in the future with initial values equal to the PIG for each state-action pair:

$$Q_0(a, s) := I_G(a, s) \quad (14)$$

Then propagating backwards in time, we maintain a




Fig. 4. Coordinating exploration using predicted information gain. The average missing information is plotted over exploration time for greedy and value-iterated (VI) maximization of PIG. The standard control strategies and the VI+ positive control are also depicted. Standard errors are plotted as dotted lines above and below learning curves. ($n=200$)

running total of estimated future value by:

$$V_\tau(s) := \max_a Q_\tau(a, s)$$

$$Q_{\tau-1}(a, s) := I_G(a, s) + \gamma \sum_{s' \in \mathcal{S}} \hat{\Theta}_{s,a,s'} \cdot V_\tau(s') \quad (15)$$

Here, $0 \leq \gamma \leq 1$, is a discount factor reducing the value of gains obtained further in the future. When $\gamma < 1$, backward propagation can be continued until convergence. Alternatively, it can simply be executed for a predefined number of steps. Choosing the latter with $\gamma = 1$, we construct a behavioral policy (PIG(VI)) for an agent that coordinates its actions under VI towards maximizing PIG:

$$a_{\text{PIG(VI)}} := \arg \max_a Q_{-10}(a, s); \quad (16)$$

Comparing the learning curves in Fig. 4 for PIG(VI) and PIG(greedy) in the three classes of worlds we find that coordination of actions yielded the greatest learning gains in Mazes, with moderate gains also seen in 1-2-3 Worlds. In Dense Worlds PIG(VI), like PIG(greedy) and random action, essentially reached maximal learning performance. Along with the results for the embodiment index above, these results support the hypothesis that worlds with high embodiment constraint require agents to coordinate their actions over several time steps to achieve efficient exploration.

Convergence and optimality of the VI algorithm can be guaranteed [8], but only if the utility function is stationary and the true world structure is known. To assess the impairment resulting from the use of the internal model in VI (15), we constructed a second positive control, PIG(VI+), which is given the true CMC kernel Θ for use during coordinated maximization of PIG under VI. Under this strategy, Θ is used only to coordinate the selection of actions and is not incorporated into the Bayesian estimate or the PIG utility function. Comparing the PIG(VI) agent to the PIG(VI+) control, we find that they only differ in Mazes, and this difference

is relatively small compared to the gains made over random or greedy behaviors (Fig. 4). Altogether these results suggest that PIG(VI) may be an effective strategy employable by embodied agents for coordinating explorative actions towards learning.

From the results so far the picture emerges that the three classes of environments offer very different challenges for the exploring agent. Dense Worlds are easy to explore. Mazes require policies that coordinate actions over time but exhibit little sensitivity to the particular choice in utility function. 1-2-3 Worlds also require coordination of actions over time, though to a lesser extent than Mazes. Unlike in Mazes, however, agents in 1-2-3 Worlds strongly benefit from the information-theoretically derived utility function PIG.

3.4 Structural features of the three worlds

We next asked how structural differences in the three classes of environments correlated with the above differences in exploration performance. In particular we considered two structural features of the worlds, their tendency to draw agents into a biased distribution over states and how tightly an action controls the future states of the agent.

State bias: To assess how strongly a world biases the state distribution of agents we consider the equilibrium distribution under an undirected action policy, random action. The equilibrium distribution Ψ is the limit distribution over states after many time steps. To quantify the bias of this distribution, we compute a *structure index* (SI) as the relative difference between its entropy $H(\Psi)$ and the entropy of the uniform distribution $H(U)$:

$$SI(\Psi) := \frac{H(U) - H(\Psi)}{H(U)}$$

where:

$$H(p(s)) := - \sum_{s \in \mathcal{S}} p(s) \log_2(p(s))$$




Fig. 5. Quantifying the structure of the worlds. (a) The embodiment index, defined in Section 3.3, is plotted against the structure index for each of 200 Dense Worlds, Mazes, and 1-2-3 Worlds. (b) The average controllability, as measured by the mutual information between an action and a future state, is plotted as a function of the number of time steps the state lies in the future ($n=200$). The error bars depict standard deviations.

The structure index values for 200 worlds in each class of environment are plotted against the embodiment index (defined in section 3.3) in Fig. 5a. As depicted, the embodiment index correlates strongly with the structure index. Thus, the state bias seems to represent a significant challenge embodied agents face during exploration.

Controllability: To measure the capacity for an agent to control its state trajectory we computed the mutual information between a random action and a future state:

$$MI[A_0, S_t | s_0] = \sum_{a_0 \in \mathcal{A}, s_t \in \mathcal{S}} p(a_0, s_t | s_0) \log_2 \left(\frac{p(s_t | a_0, s_0)}{p(s_t | s_0)} \right)$$

As shown in Fig. 5b, an action in a Maze or 1-2-3 Worlds has significantly more impact on future states than an action in Dense Worlds. Controllability is required for effective coordination of actions, such as under PIG(VI). In Mazes, where actions can significantly effect states far into the future, agents yielded the largest gains from coordinated actions. However, controllability, while necessary, is not sufficient for coordinated actions to have the potential of improving learning. For example, a non-ergodic world might have high controllability but not allow an embodied agent to ever reach a large isolated set of states, regardless of whether it coordinated its actions or not. In such a world, an unembodied agent could reach the isolated states and thereby gain a learning opportunity inaccessible to any embodied agent.

3.5 Comparison to previous explorative strategies

While exploration in the RL literature has largely focused on its role in reward acquisition, many of the principles developed to induce exploration can be implemented in our framework. In this section, we compare these various methods to PIG(VI) under our learning objective.

Random action is perhaps the most common exploration strategy used in RL. As we have already seen in Fig. 4, random action is only efficient for exploring Dense Worlds. In addition to undirected random action, the following directed exploration strategies have been developed in the RL literature. The learning curves of the various strategies are plotted in Fig. 6.

Least Taken Action (LTA): Under LTA, an agent will always choose the action that has been performed least often in the current state [7], [58], [61]. Like random action, LTA yields uniform sampling of actions in each state. Consistently, LTA fails to significantly improve on the learning rates seen under random action ($p > 0.001$ for all three environments).

Counter-Based Exploration (CB): Whereas LTA actively samples actions uniformly, CB attempts to induce a uniform sampling across states. To do this, it maintains a count of the occurrences of each state, and chooses its action to minimize the expected count of the resultant state [71]. As shown in Fig. 6, CB performs even worse than random action in Dense Worlds and 1-2-3 Worlds ($p < 0.001$). It does outperform random actions in Mazes but falls far short of the performance seen by PIG(VI) ($p < 0.001$).

Q-learning on Posterior Expected Information Gain (PEIG(Q)): Stork *et al.* [67] developed a utility function U_{Storck} to measure past changes in the internal model, which they used to guide exploration under a Q-learning algorithm [69]. Let τ be the most recent time step in the past over which the internal model for the transition distribution $\Theta_{a,s,:}$ changed:

$$\tau := \max\{t | s(t) = s, a(t) = a, t < |\vec{d}|\}$$

Then, considering the internal model before and after this time step (Θ^τ and $\Theta^{\tau+1}$ respectively), and the




Fig. 6. Comparison to previous exploration strategies. The average missing information is plotted over time for PIG(VI) agents along with three exploration strategies from the literature: least taken action (LTA) [7], [58], [61], counter-based (CB) [71], and Q-Learning on posterior expected information gain (PEIG(Q)) [67]. The standard control strategies are also shown. Standard errors are plotted as dotted lines above and below learning curves. ($n=200$)

data collected up to this point $\vec{d}^{\tau+1}$, the utility function defined by Storck *et al.* is:

$$U_{\text{Storck}} := D_{\text{KL}}[\hat{\Theta}_{a,s,:}^{\tau+1} \parallel \hat{\Theta}_{a,s,:}^{\tau}] \quad (17)$$

Note, both $\hat{\Theta}^{\tau}$ and $\hat{\Theta}^{\tau+1}$ are internal models previously (or currently) held by the agent. In the following derivation, we demonstrate that U_{Storck} is equivalent to the posterior expected information gained (PEIG).

$$\begin{aligned} \text{PEIG}(a, s) &:= E_{\Theta|\vec{d}^{\tau+1}} \left[I_M(\Theta \parallel \hat{\Theta}^{\tau}) - I_M(\Theta \parallel \hat{\Theta}^{\tau+1}) \right] \\ &= E_{\Theta|\vec{d}^{\tau+1}} \left[\sum_{s'} \Theta_{a,s,s'} \log_2 \left(\frac{\hat{\Theta}_{a,s,s'}^{\tau+1}}{\hat{\Theta}_{a,s,s'}^{\tau}} \right) \right] \\ &= \sum_{s'} E_{\Theta|\vec{d}^{\tau+1}} [\Theta_{a,s,s'}] \log_2 \left(\frac{\hat{\Theta}_{a,s,s'}^{\tau+1}}{\hat{\Theta}_{a,s,s'}^{\tau}} \right) \\ &= \sum_{s'} \hat{\Theta}_{a,s,s'}^{\tau+1} \log_2 \left(\frac{\hat{\Theta}_{a,s,s'}^{\tau+1}}{\hat{\Theta}_{a,s,s'}^{\tau}} \right) \\ &= D_{\text{KL}}[\hat{\Theta}_{a,s,:}^{\tau+1} \parallel \hat{\Theta}_{a,s,:}^{\tau}] \end{aligned} \quad (18)$$

Thus, PEIG is a posterior analogue to our PIG utility function. Q-learning is a model-free approach to maximizing long-term gains of a utility function [69]. Following Storck *et al.*, we tested the combination of PEIG and Q-learning (PEIG(Q)) in our test environments. Surprisingly, PEIG(Q) performs even worse, at least initially, than random action in all three environments ($p < 0.001$ for CMCs and 1-2-3 Worlds, $p > 0.001$ for Mazes). As such, it fails to yield the learning performance seen by PIG(VI) in Mazes in 1-2-3 Worlds.

Altogether, Fig. 6 demonstrates that PIG(VI) outperforms the previous explorative strategies at learning structured worlds. To further compare the principles of PIG(VI) and PEIG(Q), we introduce two cross-over strategies that borrow from each of them. The first is PIG(Q) which applies Q-learning to the PIG utility function. The learning performance of PIG(Q) is similar to PEIG(Q), falling short of PIG(VI) (Fig. S3). This suggests

that Q-learning is ineffective at coordinating actions during exploration. The second cross-over strategy is PEIG(VI) which applies the VI algorithm to Storck *et al.*'s utility function. PEIG(VI) matched PIG(VI) in Mazes ($p > 0.001$) but not 1-2-3 Worlds ($p < 0.001$), suggesting that the posterior information gain is a reasonable predictor for future information gain under a Dirichlet prior but not a Discrete prior.

3.6 Comparison to utility functions from Psychology

Inspired by independent findings in the field of Psychology that PIG can describe human behavior during hypothesis testing, we investigated two other measures also developed in this context [44], [46]. Like PIG, both are measures of the difference between the current and hypothetical future internal models:

Predicted mode change (PMC) predicts the height difference between the modes of the current and future internal models [6], [44]:

$$\text{PMC}(a, s) = \sum_{s^*} \hat{\Theta}_{s,a,s^*} \left[\max_{s'} \hat{\Theta}_{a,s,s'}^{a,s,s^*} - \max_{s'} \hat{\Theta}_{a,s,s'} \right] \quad (19)$$

Predicted L1 change (PLC) predicts the average L1 distance between the current and future internal models [35]:

$$\text{PLC}(a, s) = \sum_{s^*} \hat{\Theta}_{s,a,s^*} \left[\frac{1}{N} \sum_{s'} \left| \hat{\Theta}_{a,s,s'}^{a,s,s^*} - \hat{\Theta}_{a,s,s'} \right| \right] \quad (20)$$

Note, PMC and PLC differ from PIG in the norm used to quantify differences between CMC kernels. Considering an arbitrary norm d , the claim analogous to Theorem 2 would be:

$$\begin{aligned} &\sum_{s^*} \hat{\Theta}_{a,s,s^*}^{\text{current}} d(\hat{\Theta}^{\text{future}} \parallel \hat{\Theta}^{\text{current}}) \\ &= E_{s^*, \Theta|\vec{d}} \left[d(\Theta \parallel \hat{\Theta}^{\text{current}}) - d(\Theta \parallel \hat{\Theta}^{\text{future}}) \right] \end{aligned}$$




Fig. 7. Comparison between utility functions. The average missing information is plotted over time for agents that employ VI to maximize long-term gains in the three objective function, PIG, PMC, or PLC. The standard control strategies are also shown. Standard errors are plotted as dotted lines above and below learning curves. ($n=200$)

This claim states that the expected difference between the current and future internal model equals the expected change in difference with respect to the ground truth. While this claim holds when d is the norm used in PIG (Theorem 2), it does not generally hold for either of the norms used in PMC or PLC.

To our knowledge, neither PIG, PMC nor PLC have previously been applied to sequences of observations or to embodied action perception loops. We tested agents that attempt to maximize PMC or PLC using VI. As Fig. 7 reveals, PIG(VI) proved again to be the best performer overall. In particular, PIG(VI) significantly outperforms PMC(VI) in all three environments, and PLC(VI) in 1-2-3 Worlds ($p < 0.001$). Nevertheless, PMC and PLC achieved significant improvements over the baseline control in Mazes and 1-2-3 Worlds, highlighting the benefit of value iteration across different utility functions. Interestingly, when performance was measured by an L1 distance instead of missing information, PIG(VI) still outperformed PMC(VI) and PLC(VI) in 1-2-3 Worlds (Fig. S4).

3.7 Generalized utility of exploration

From a behavioral perspective, learning represents a fundamental and primary drive [2], [37]. The evolutionary advantage of such an exploratory drive likely rests on the general utility of the acquired internal model [33], [50], [51], [55], [56]. To test this, we assessed the ability of the agents to use their internal models, derived through exploration, to accomplish an array of goal-directed tasks. We consider two groups of tasks: navigation and reward acquisition.

Navigation: Starting at any given state, the agent has to navigate to any given target state with the minimal number of steps.

Reward Acquisition: For every starting state, the agent has to gather as much reward as possible over 100 time steps. Reward values are drawn from a standard normal distribution and randomly assigned to every state in the

CMC. Each agent is tested in ten randomly generated reward structures.

At several time points during exploration, the agent is stopped and its internal models assessed for general utility. For each task, we next derive the behavioral policy that optimizes performance under the internal model. The derived policy is then tested in the world (i.e. under the true CMC kernel), and the expected path length or acquired reward for that policy is determined. As a positive control, we also derive an objective optimal policy that maximizes the realized performance for the true CMC kernel. The difference in realized performance between the subjective and objective policies is used as a measure of navigational loss or reward loss. High navigational loss means the agents policy took many more time steps to reach the target state than the optimal policy. High reward loss means the agents policy yielded significantly fewer rewards than the optimal policy.

Fig. 8 depicts the average rank in navigational and reward loss for the different explorative strategy. Significance bounds ($p = 0.001$) around PIG(VI) were determined by post-hoc analysis of Friedman’s test [25]. In all environments, for both navigation and reward acquisition, PIG(VI) always grouped with the top performers ($p > 0.001$), excepting positive controls. PIG(VI) was the only strategy to do so. Thus, the explorative strategy that optimized learning under the missing information objective function gave the agent an advantage in a range of independent tasks.

4 DISCUSSION

In this manuscript we introduced a parsimonious mathematical framework for studying learning-driven exploration by embodied agents based on information theory, Bayesian inference and controllable Markov chains (CMCs). We compared agents that utilized different exploration strategies towards optimizing learning. To understand how learning performance depends on the structure of the world, three classes of environments




Fig. 8. Demonstration of generalized utility. For each world ($n=200$), explorative strategies are ranked for average navigational loss (averaged across N start states and N target states) and average reward loss (averaged across N start states and 10 randomly generated reward distributions). The average ranks are plotted with standard deviations. Strategies lying outside the pair of solid green lines differ significantly from PIG(VI) in navigational loss. Strategies lying outside the pair of solid blue lines differ significantly from PIG(VI) in reward loss ($p < 0.0001$). The different utility functions and heuristics are distinguished by color: PIG(green), PEIG(magenta), PMC(blue), PLC(cyan), LTA(orange), CB(yellow). The different coordination methods are distinguished by symbol: Greedy(squares), VI(circles), VI+(diamonds), Heuristic Strategies(asterisks). The two standard controls are depicted as follows: Unembodied(black), Random(red).

were considered that challenge the learning agent in different ways. We found that fast learning could be achieved by an exploration strategy that coordinated actions towards long-term maximization of predicted information gain (PIG).

4.1 Potential limitations to our approach

The optimality of the Bayesian Estimate (Theorem 1) and the accuracy of PIG (Theorem 2) both require a prior distribution on the transition distributions. For biological agents, such priors could have been learned from earlier exploration of related environments, or may represent hardwired beliefs optimized by evolutionary pressures. As another possibility, an agent could attempt to simultaneously learn a prior while exploring its environment. Indeed, additional results (Fig. S5) show that the maximum-likelihood estimation of the concentration parameter for Dense Worlds and Mazes enables exploration that quickly matches the performance of agents given accurate priors. Nevertheless, biological agents may not always have access to an accurate prior for an environment. For such cases, future work is required to understand exploration under false priors and how they could yield sub-optimal but perhaps biologically realistic exploratory behaviors.

As another potential limitation of our approach, the VI algorithm is only optimal for dynamic processes with known stationary transition probabilities and stationary utilities [8]. In contrast, any utility function, including PIG, that attempts to capture the progress in learning of an agent will necessarily change over time. This caveat may be partially alleviated by the fact that PIG changes only for the sampled distributions. Furthermore, PIG decreases in a monotonic fashion (see Fig. 2) which could

potentially be captured by the discount factor of VI. Interesting future work may lie in accounting for the effect of such monotonic decreases in estimates of future learning gains.

In addition, the learning agent does not have access to the true transition distributions for performing VI and has to rely instead on its evolving internal model. The impairment caused by this reliance on the internal model was directly assessed with a positive control PIG(VI+). A comparison of PIG(VI) against this control (Fig. 4) shows performance impairment only in Mazes and it is rather small compared to the improvements offered by VI.

Finally, it might be argued that the use of missing information as a measure of learning unfairly advantaged the PIG utility function. Interestingly, however, PIG under VI was not only the fastest learner, but also demonstrated the greatest capacity for accomplishing goal-directed tasks. Furthermore, it even outperformed other strategies, including PLC(VI), under an L1 objective function (Fig. S4).

4.2 Related work in Reinforcement Learning

CMCs are closely related to the more commonly studied Markov Decision Processes (MDPs) used in Reinforcement Learning. MDPs differ from CMCs in that they explicitly include a stationary reward function associated with each transition [19], [69]. RL research of exploration usually focusses on its role in balancing exploitative behaviors during reward maximization. Several methods for inducing exploratory behavior in RL agents have been developed. Heuristic strategies such as random action, least taken action, and counter-based algorithms are commonly employed in the RL literature. While such strategies may be useful in RL, our results show that they

are inefficient for learning the dynamics of structured worlds.

In contrast to these heuristic strategies for exploration, several principled approaches have been proposed for inducing exploratory actions to maximize rewards. For example, the BEETLE algorithm models reward as a partially observable MDP and derives an analytic solution to optimize rewards [53]. Similarly, the BOSS approach maintains a posterior distribution over MDPs from which it periodically samples for selecting actions that maximize reward gains “optimistically” from the samples [3]. These strategies focus exclusively on extrinsically motivated exploration and do not address exploration driven by learning for its own sake.

Finally, several studies have investigated intrinsically motivated learning under the RL framework. For example, Singh et al. [63] have demonstrated that RL guided by saliency, an intrinsic motivation derived from changes in stimulus intensity, can promote the learning of reusable skills. As mentioned previously, Storck et al. introduced the combination of Q-learning and PEIG as an intrinsic motivator of learning [67]. In their study, PEIG(Q) outperformed random action only over long time scales. At shorter time scales, random action performed better. Interestingly, we found exactly the same trend, initially slow learning with eventual catching-up, when we applied PEIG(Q) to exploration in our test environments (Fig. 6).

4.3 Related work in Psychology

In the Psychology literature, PIG, as well as PMC and PLC, were directly introduced as measures of the expected difference between a current and future belief [6], [35], [44], [46]. Here, in contrast, we derived PIG, using Bayesian inference, from the expected change in missing information with respect to a ground truth (Theorem 2). Analogous theorems do not hold for PMC or PLC. For example, the expected change in L1 distance between an internal model and the true structure is not equivalent to the expected L1 distance between successive internal models. This might explain why PIG(VI) outperformed PLC(VI) even under an L1 measure of learning (Fig. S4).

We applied the PIG principle to the learning of a full model of the world. The Psychology literature, in contrast, focusses on specific questions (hypothesis testing) regarding the data. In addition, this prior literature has not considered sequences of actions or embodied sensor-motor loops.

It has been shown that human behavior during hypothesis testing can be explained by a model that maximizes PIG [44], [46]. This suggests that the PIG information-theoretic measure may have biological significance. The behavioral studies, however, could not distinguish between the different utility functions (PIG, PMC and PLC) in their ability to explain human behavior [44]. Perhaps our finding that 1-2-3 Worlds give rise to large differences between the three utility functions

can help identify new behavioral tasks for disambiguating the role of these measures in human behavior.

Itti and Baldi recently developed an information theoretic measure closely related to PEIG for modeling bottom-up visual saliency and predicting gaze attention [5], [29], [30]. In their model, a Bayesian learner maintains a probabilistic belief structure over the low-level features of a video. Attention is believed to be attracted to locations in the visual scene that exhibit high Surprise. Like PEIG, Surprise quantifies changes in posterior beliefs by a summed Kullback-Leibler divergence. Several potential extensions of this work are suggested by our results. First, it may be useful to model the active nature of data acquisition during visual scene analysis. In Itti and Baldi’s model, all features are updated for all location of the visual scene regardless of current gaze location or gaze trajectory. Differences in acuity between the fovea and periphery however suggest that gaze location will have a significant effect on which low-level features can be transmitted by the retina [74]. Second, our comparison between the PIG and PEIG utility functions (Figs. 6 and S3) suggests that predicting where future change might occur, may be more efficient than focusing attention only on those locations where change has occurred in the past. A model that anticipates Surprise, as PIG anticipates information gain, may be better able to explain some aspects of human attention. For example, if a moving subject disappears behind an obscuring object, viewers may anticipate the reemergence of the subject and attend the far edge of the obscurer. Incorporating these insights into new models of visual saliency and attention could be an interesting course of future research.

4.4 Information-theoretic models of behavior

The field of behavioral modeling has recently seen increased utilization of information-theoretic concepts. These approaches can be grouped under three guiding principles. The first group uses information theory to quantify the complexity of a behavioral policy, with high complexity generally considered undesirable. Tishby and Polani for example, considered RL maximization of rewards under such complexity constraints [73]. While we did not consider complexity constraints on our behavioral strategies in the current work, it may be an interesting topic for future studies.

The second common principle seeks to maximize *predictive information* [4], [66], [72] (not to be confused with predicted information gain, PIG). Predictive information, which has also been termed *excess entropy* [14], estimates the amount of information a known variable (or past variable) contains regarding an unknown (or future) variable. For example, in simulated robots, Ay et al. demonstrated that complex and interesting behaviors can emerge by choosing control parameters that maximize the predictive information between successive sensory inputs [4]. The information bottleneck approach introduced by Tishby et al. [72] combines predictive

information and complexity constraints, maximizing the information between a compressed internal variable and future state progression subject to a constraint on the complexity of generating the internal variable from sensory inputs. Recently, Still extended the information bottleneck method to incorporate actions [66].

Both Ay *et al.* and Still describe the behaviors that result from their models as exploratory. Their objective of high predictive information selects actions such that the resulting sensory input changes often but in a predictable way. We therefore call this form of exploration stimulation-driven. Predictive information can only be high when the sensory feedback can be predicted, and thus stimulation-driven exploration relies on an accurate internal model. In contrast, the learning objective we introduced here drives actions most strongly when the internal model can be improved and this drive weakens as it becomes more accurate. Thus, learning-driven and stimulation-driven exploration contrast each other while being very interdependent. Indeed, a simple additive combination of the two objectives may naturally lead to a smooth transitioning between the two types of exploration, directed by the expected accuracy of the internal model. In the next section we suggest a correspondence of these two computational principles of exploration with two distinct modes of behavior distinguished in psychology and behavioral research.

Finally, the Free-Energy (FE) hypothesis introduced by Friston proposes that the minimization of free-energy, an information-theoretic bound on surprise, offers a unified variational principle for governing both the learning of an internal model as well as actions [17]. Friston notes that under this principle agents should act to minimize the number of states they visit. This stands in stark contrast to both learning-driven and stimulation-driven exploration. During learning-driven exploration, an agent will seek out novel states where missing information is high. During stimulation-driven exploration, an agent will actively seek to maintain high variation in its sensory inputs. Nevertheless, as Friston argues, reduced state entropy may be valuable in dangerous environments where few states permit survival. The balance between cautionary and exploratory behaviors would be an interesting topic for future research.

4.5 Towards a general theory of exploration

With the work of Berlyne [11], Psychologists began to dissect the complex domains of behavior and motivation that comprise exploration. A distinction between play (or diversive exploration) and investigation (or specific exploration) grew out of two competing theories of exploration. As reviewed by Hutt [27], "curiosity"-theory proposed that exploration is a consummatory response to curiosity-inducing stimuli [9], [42]. In contrast, "boredom"-theory held that exploration was an instrumental response for stimulus change [20], [43]. To ameliorate this opposition, Hutt suggested that the two

theories may be capturing distinct behavioral modes, with "curiosity"-theory underlying investigatory exploration and "boredom"-theory underlying play. In children, exploration often occurs in two stages, inspection to understand what is perceived, followed by play to maintain changing stimulation [28]. These distinctions nicely correspond to the differences between our approach and the predictive information approach of Ay *et al.* [4] and Still [66]. In particular, we hypothesize that our approach, which emphasizes the acquisition of information, corresponds to curiosity-driven investigation. In contrast, we propose that predictive information a la Ay *et al.* and Still, which rehearses the internal model in a wide range, may correspond with play. Further, the proposed method of additively combining these two principles (Section 4.4), may naturally capture the transition between investigation and play seen in children during exploration.

Even in the domain of curiosity-driven exploration, there are many varied theories [37]. Early theories viewed curiosity as a drive to maintain a specific level of arousal. These were followed by theories interpreting curiosity as a response to intermediate levels of incongruence between expectations and perceptions, and later by theories interpreting curiosity as a motivation to master one's environment. Loewenstein developed an Information Gap Theory and suggested that curiosity is an aversive reaction to missing information [37]. More recently, Silvia proposed that curiosity is composed of two appraisal components, complexity and comprehensibility. For Silvia complexity is broadly defined, and includes novelty, ambiguity, obscurity, mystery, etc. Comprehensibility appraises whether something can be understood. It is interesting how well these two appraisals match information-theoretic concepts, complexity being captured by entropy, and comprehensibility by information gain [49]. Indeed, predicted information gain might be able to explain the dual appraisals of curiosity-driven exploration proposed by Silvia. PIG is bounded by entropy and thus high values require high complexity. At the same time, PIG equals the expected decrease in missing information and thus may be equivalent to expected comprehensibility.

All told, our results add to a bigger picture of exploration in which the theories for its different aspects fit together like pieces of a puzzle. This invites future work for integrating these pieces into a more comprehensive theory of exploration and ultimately of autonomous behavior.

ACKNOWLEDGEMENTS

The authors wish to thank Nihat Ay, Susanne Still, Jim Crutchfield, Reza Moazzezi, and the Redwood Center for Theoretical Neuroscience for useful discussions. D.Y. Little is supported by the Redwood Center for Theoretical and Computational Neuroscience Graduate Student Endowment. This work was supported in part by the National Science Foundation grant CNS-0855272.

REFERENCES

- [1] P. Abbeel and A. Ng, "Exploration and apprenticeship learning in reinforcement learning," in *Proceedings of the 22nd international conference on Machine learning*. ACM, 2005, pp. 1–8.
- [2] J. Archer and L. Birke, *Exploration in animals and humans*. Van Nostrand Reinhold (UK) Co. Ltd., 1983.
- [3] J. Asmuth, L. Li, M. Littman, A. Nouri, and D. Wingate, "A bayesian sampling approach to exploration in reinforcement learning," in *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2009, pp. 19–26.
- [4] N. Ay, N. Bertschinger, R. Der, F. Güttler, and E. Olbrich, "Predictive information and explorative behavior of autonomous robots," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 63, no. 3, pp. 329–339, 2008.
- [5] P. F. Baldi and L. Itti, "Of bits and wows: A bayesian theory of surprise with applications to attention," *Neural Networks*, vol. 23, no. 5, pp. 649–666, Jun 2010.
- [6] J. Baron, "Rationality and intelligence." 1985.
- [7] A. Barto and S. Singh, "On the computational economics of reinforcement learning," in *Connectionist Models: Proceedings of the 1990 Summer School*. Morgan Kaufmann. Citeseer, 1990.
- [8] R. E. Bellman, *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957.
- [9] D. Berlyne, "Novelty and curiosity as determinants of exploratory behaviour1," *British Journal of Psychology. General Section*, vol. 41, no. 1-2, pp. 68–80, 1950.
- [10] —, "Conflict, arousal, and curiosity." 1960.
- [11] —, "Curiosity and exploration," *Science*, vol. 153, no. 3731, p. 25, 1966.
- [12] J. Byers, *Animal play: evolutionary, comparative, and ecological perspectives*. Cambridge Univ Pr, 1998.
- [13] T. Cover and J. Thomas, *Elements of information theory*. Wiley Online Library, 1991, vol. 6.
- [14] J. Crutchfield and D. Feldman, "Regularities unseen, randomness observed: levels of entropy convergence," *Chaos (Woodbury, NY)*, vol. 13, no. 1, p. 25, 2003.
- [15] J. Crutchfield and B. McNamara, "Equations of motion from a data series," *Complex systems*, vol. 1, no. 3, pp. 417–452, 1987.
- [16] S. Dura-Bernal, T. Wennekers, and S. Denham, "Modelling object perception in cortex: Hierarchical bayesian networks and belief propagation," in *Information Sciences and Systems (CISS), 2011 45th Annual Conference on*, march 2011, pp. 1–6.
- [17] K. Friston, "The free-energy principle: a rough guide to the brain?" *Trends in cognitive sciences*, vol. 13, no. 7, pp. 293–301, 2009.
- [18] D. George and J. Hawkins, "A hierarchical bayesian model of invariant pattern recognition in the visual cortex," in *Neural Networks, 2005. IJCNN '05. Proceedings. 2005 IEEE International Joint Conference on*, vol. 3, july-4 aug. 2005, pp. 1812–1817 vol. 3.
- [19] H. Gimbert, "Pure stationary optimal strategies in markov decision processes," *STACS 2007*, pp. 200–211, 2007.
- [20] M. Glanzer, "Curiosity, exploratory drive, and stimulus satiation," *Psychological Bulletin*, vol. 55, no. 5, p. 302, 1958.
- [21] S. Glickman and R. Sroges, "Curiosity in zoo animals," *Behaviour*, pp. 151–188, 1966.
- [22] G. Gordon, D. Kaplan, B. Lankow, D. Little, J. Sherwin, B. Suter, and L. Thaler, "Toward an integrated approach to perception and action: Conference report and future directions," *Frontiers in Systems Neuroscience*, vol. 5, 2011.
- [23] R. W. Guillery, "Anatomical pathways that link perception and action," *Prog. Brain Res.*, vol. 149, pp. 235–256, 2005.
- [24] R. W. Guillery and S. M. Sherman, "Branched thalamic afferents: what are the messages that they relay to the cortex?" *Brain Res Rev*, vol. 66, pp. 205–219, Jan 2011.
- [25] Y. Hochberg and A. Tamhane, *Multiple comparison procedures*. Wiley Online Library, 1987, vol. 82.
- [26] S. Holmes, "The selection of random movements as a factor in phototaxis," *Journal of Comparative Neurology and Psychology*, vol. 15, no. 2, pp. 98–112, 1905.
- [27] C. Hutt, "Specific and diverse exploration1," *Advances in child development and behavior*, vol. 5, p. 119, 1970.
- [28] C. Hutt and R. Bhavnani, "Predictions from play." *Nature*, 1972.
- [29] L. Itti and P. Baldi, "Bayesian surprise attracts human attention," *Advances in neural information processing systems*, vol. 18, p. 547, 2006.
- [30] L. Itti and P. F. Baldi, "Bayesian surprise attracts human attention," *Vision Research*, vol. 49, no. 10, pp. 1295–1306, May 2009.
- [31] H. S. Jennings, *Behavior of the lower organisms*. New York, The Columbia university press, The Macmillan company, agents, 1906.
- [32] L. Kaelbling, M. Littman, and A. Moore, "Reinforcement learning: A survey," *Arxiv preprint cs/9605103*, 1996.
- [33] R. Kaplan and S. Kaplan, "Cognition and environment: functioning of an uncertain world. ulrich's bookstore," *Ann Arbor*, 1983.
- [34] M. Kawato and K. Samejima, "Efficient reinforcement learning: computational theories, neuroscience and robotics," *Current opinion in neurobiology*, vol. 17, no. 2, pp. 205–212, 2007.
- [35] J. Klayman and Y. Ha, "Confirmation, disconfirmation, and information in hypothesis testing," *Psychological review*, vol. 94, no. 2, p. 211, 1987.
- [36] M. Lewicki, "Efficient coding of natural sounds," *nature neuroscience*, vol. 5, no. 4, pp. 356–363, 2002.
- [37] G. Loewenstein, "The psychology of curiosity: A review and reinterpretation," *psychological Bulletin*, vol. 116, no. 1, p. 75, 1994.
- [38] D. MacKay and L. Peto, "A hierarchical dirichlet language model," *Natural language engineering*, vol. 1, no. 3, pp. 1–19, 1995.
- [39] C. Manning, P. Raghavan, and H. Schütze, "Introduction to information retrieval," 2008.
- [40] T. Matsuzawa and M. Tanaka, *Cognitive development in chimpanzees*. Springer Verlag, 2006.
- [41] V. Mikheev and O. Andreev, "Two-phase exploration of a novel environment in the guppy, poecilia reticulata," *Journal of fish biology*, vol. 42, no. 3, pp. 375–383, 1993.
- [42] K. Montgomery, "Exploratory behavior as a function of" similarity" of stimulus situation," *Journal of Comparative and Physiological Psychology*, vol. 46, no. 2, p. 129, 1953.
- [43] A. Myers and N. Miller, "Failure to find a learned drive based on hunger; evidence for learning motivated by" exploration." *Journal of Comparative and Physiological Psychology*, vol. 47, no. 6, p. 428, 1954.
- [44] J. Nelson, "Finding useful questions: on bayesian diagnosticity, probability, impact, and information gain," *Psychological Review*, vol. 112, no. 4, p. 979, 2005.
- [45] A. Noë, *Action in perception*. the MIT Press, 2004.
- [46] M. Oaksford and N. Chater, "A rational analysis of the selection task as optimal data selection," *Psychological Review*, vol. 101, no. 4, p. 608, 1994.
- [47] B. Olshausen *et al.*, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.
- [48] J. O'Regan and A. Noë, "A sensorimotor account of vision and visual consciousness," *Behavioral and brain sciences*, vol. 24, no. 5, pp. 939–972, 2001.
- [49] E. Pfaffelhuber, "Learning and information theory," *International Journal of Neuroscience*, vol. 3, no. 2, pp. 83–88, 1972.
- [50] W. Pisula, "Costs and benefits of curiosity: The adaptive value of exploratory behavior," *Polish Psychological Bulletin*, 2003.
- [51] —, "Play and exploration in animalsa comparative analysis," *Polish Psychological Bulletin*, vol. 39, no. 2, pp. 104–107, 2008.
- [52] —, *Curiosity and information seeking in animal and human behavior*. Brown Walker Pr, 2009.
- [53] P. Poupart, N. Vlassis, J. Hoey, and K. Regan, "An analytic solution to discrete bayesian reinforcement learning," in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 697–704.
- [54] M. Rehn and F. Sommer, "A network that uses few active neurons to code visual input predicts the diverse shapes of cortical receptive fields," *Journal of Computational Neuroscience*, vol. 22, no. 2, pp. 135–146, 2007.
- [55] M. J. Renner, "Learning during exploration: the role of behavioral topography during exploration in determining subsequent adaptive behavior," *Int J Comp Psychol*, vol. 2, no. 1, p. 4356, 1988.
- [56] M. Renner, "Neglected aspects of exploratory and investigatory behavior," *Psychobiology*, 1990.
- [57] P. Rochat, "Object manipulation and exploration in 2-to 5-month-old infants," *Developmental Psychology*, vol. 25, no. 6, p. 871, 1989.
- [58] M. Sato, K. Abe, and H. Takeda, "Learning control of finite markov chains with an explicit trade-off between estimation and control," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 18, no. 5, pp. 677–684, 1988.
- [59] T. Serre, L. Wolf, S. Bileschi, M. Riesenhuber, and T. Poggio, "Robust object recognition with cortex-like mechanisms," *IEEE transactions on pattern analysis and machine intelligence*, pp. 411–426, 2007.

- [60] J. Shinskey and Y. Munakata, "Something old, something new: a developmental transition from familiarity to novelty preferences with hidden objects," *Developmental science*, vol. 13, no. 2, pp. 378–384, 2010.
- [61] B. Si, J. Herrmann, and K. Pawelzik, "Gain-based exploration: From multi-armed bandits to partially observable environments," 2007.
- [62] P. Silvia, "What is interesting? exploring the appraisal structure of interest," *Emotion*, vol. 5, no. 1, p. 89, 2005.
- [63] S. Singh, R. Lewis, A. Barto, and J. Sorg, "Intrinsically motivated reinforcement learning: An evolutionary perspective," *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 2, pp. 70–82, 2010.
- [64] A. Stewart, J. Cachat, K. Wong, N. Wu, L. Grossman, C. Suci, J. Goodspeed, M. Elegante, B. Bartels, S. Elkhayat *et al.*, "Phenotyping of zebrafish homebase behaviors in novelty-based tests," *Neuromethods*, vol. 51, pp. 143–156, 2011.
- [65] A. Stewart, S. Gaikwad, E. Kyzar, J. Green, A. Roth, and A. Kaluff, "Modeling anxiety using adult zebrafish: A conceptual review," *Neuropharmacology*, 2011.
- [66] S. Still, "Information-theoretic approach to interactive learning," *EPL (Europhysics Letters)*, vol. 85, p. 28005, 2009.
- [67] J. Storck, S. Hochreiter, and J. Schmidhuber, "Reinforcement driven information acquisition in non-deterministic environments," in *ICANN'95*. Citeseer, 1995.
- [68] M. Stowe, T. Bugnyar, M. Loretto, C. Schloegl, F. Range, and K. Kotrschal, "Novel object exploration in ravens (*corvus corax*): effects of social relationships," *Behavioural processes*, vol. 73, no. 1, pp. 68–75, 2006.
- [69] R. Sutton and A. Barto, *Reinforcement learning*. MIT Press, 1998, vol. 9.
- [70] R. Sutton and B. Pinette, "The learning of world models by connectionist networks," in *Proceedings of the seventh annual conference of the cognitive science society*. Citeseer, 1985, pp. 54–64.
- [71] S. Thrun, "Efficient exploration in reinforcement learning," *Technical report, School of Computer Science, Carnegie-Mellon University*, 1992.
- [72] N. Tishby, F. Pereira, and W. Bialek, "The information bottleneck method," *Proceedings of the 37th Allerton Conference on Communication, Control and Computation*, 1999.
- [73] N. Tishby and D. Polani, "Information theory of decisions and actions," *Perception-Action Cycle*, pp. 601–636, 2011.
- [74] H. Wässle, B. Boycott *et al.*, "Functional architecture of the mammalian retina," *Physiological reviews*, vol. 71, no. 2, p. 447, 1991.
- [75] A. Wüschmann, "Quantitative untersuchungen zum neugierverhalten von wirbeltieren," *Zeitschrift für Tierpsychologie*, vol. 20, no. 1, pp. 80–109, 1963.




Fig. S1. Example Dense World. Dense Worlds consist of 4 actions (separately depicted) and 10 states (depicted as nodes of the graphs). The transition probabilities associated with taking a particular action are depicted as arrows pointing from the current state to each of the possible resultant states. Arrow color depicts the likelihood of each transition.




Fig. S2. Example 1-2-3 World. 1-2-3 Worlds consist of 3 actions (separately depicted) and 20 states (depicted as nodes of the graphs). The transition probabilities associated with taking a particular action are depicted as arrows pointing from the current state to each of the possible resultant states. Arrow color depicts the likelihood of each transition. The absorbing state is depicted in gray.




Fig. S3. Comparison between different features of current and previous exploration strategies. The Average missing information is plotted over time for agents that apply either VI (circles) or Q-learning (triangles) towards maximization of either PIG (green) or PEIG (magenta). Standard control strategies are also shown. Standard errors are plotted as dotted lines above and below learning curves. ($n=200$)




Fig. S4. Comparison between utility functions under L1 objective. The average L1 distance is plotted over time for agents that coordinate actions using VI to maximize long-term gains in PIG, PMC, or PLC. Standard control strategies are also shown. Standard errors are plotted as dotted lines above and below learning curves. ($n=200$)




Fig. S5. Inferring the concentration parameter during learning. Maximum Likelihood estimation of the concentration parameter, α , is performed from data collected during exploration. A maximum concentration of $\alpha=20$ is imposed. (a,b) The mean error in inferred concentration parameter over time is plotted for Dense Worlds and Mazes. (c,d) The missing information for a PIG(VI) explorer updating its internal model using the true (green with circles) or inferred (purple with stars) over time is plotted for Dense Worlds and Mazes. Standard control explorers (with α given) have been included. Dotted lines above and below learning curves depict standard errors. Notice, even when required to infer the appropriate concentration parameter, the explorer is still able to quickly learn an accurate internal model.